arXiv:1408.0337v1 [stat.ML] 2 Aug 2014

A Bayesian estimation approach to analyze non-Gaussian data-generating processes with latent classes

Naoki Tanaka* Shohei Shimizu* Takashi Washio*

Abstract

A large amount of observational data has been accumulated in various fields in recent times, and there is a growing need to estimate the generating processes of these data. A linear non-Gaussian acyclic model (LiNGAM) based on the non-Gaussianity of external influences has been proposed to estimate the data-generating processes of variables. However, the results of the estimation can be biased if there are latent classes. In this paper, we first review LiNGAM, its extended model, as well as the estimation procedure for LiNGAM in a Bayesian framework. We then propose a new Bayesian estimation procedure that solves the problem.

1 Introduction

Several methods have recently been proposed to discover a complete causal structure, that is, all the causal directions, under the assumption that disturbance variables have non-Gaussian distributions. However, the estimation results can be biased if there are "latent classes." Latent classes are unobserved discrete variables that have more than one observed child variables. Data that has been generated from different processes are mixed in the presence of latent classes. Several methods have been proposed to estimate the causal structure in the presence of latent classes [12], but all of these are affected by local optima. Therefore, in this paper, we propose a new estimation approach that can solve this problem.

The structure of this paper is as follows. In Section 2, we briefly review the data generating process for estimating causal structure, (LiNGAM, short for Linear Non-Gaussian Acyclic Model) [11], the LiNGAM mixture model [12], existing estimation approaches to the LiNGAM mixture model [12], and the BayesLiNGAM model. In Section 3, we extend previous research dedicated to solving the problem, and test the performance of our new method through experiments with simulated data in Section 4. Section 5 concludes the paper.

*The Institute of Scientific and Industrial Research, Osaka University

2 Background

2.1 Linear non-Gaussian acyclic model (LiNGAM)

We begin by introducing the basic LiNGAM model [11]. LiNGAM is a causal model with the following four assumptions.

1. The relations between the observed variables $x_i (i = 1, \ldots, n)$ can be represented by a directed acyclic graph (DAG), as shown in Fig. 1.

2. x_i is assigned a value by a *linear function* of the values already assigned to the variables constituting its parents in the DAG as well as a "disturbance" (noise) term e_i and an optimal constant term μ_i:

$$x_i = \sum_{k(j)<k(i)} b_{ij}(x_j - \mu_j) + e_i + \mu_i \tag{1}$$

   where $k(i)$ is a causal order of x_i in the DAG (so that if there is a directed edge from x_j to x_i in the DAG, $k(j) < k(i)$), b_{ij} represents the strength of the connection between x_j and x_i in the DAG.

3. The e_i are all continuous random variables that follow *non-Gaussian* distributions with *zero means* and non-zero variances, and the e_i are mutually independent, i.e., $p(e_1, \ldots, e_n) = \prod_i p_i(e_i)$.

4. The dataset $\mathbf{D} = \{\mathbf{x}^1, \ldots, \mathbf{x}^N\}$ (each $\mathbf{x}$ contains components x_i) are observed, and each data vector $\mathbf{x}$ is generated according to the process described above with the same DAG, coefficients b_{ij}, constants μ_i, and disturbances e_i, sampled independently from the same distributions. Note that this assumption implies that there are either no unobserved (latent) confounders [10] (hidden variables) or that the x_i are causally sufficient [13].

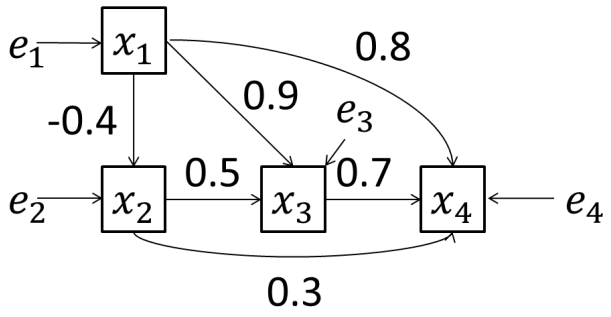

Figure 1: An example of LiNGAM described by a DAG.

2.2 LiNGAM mixture model

The LiNGAM mixture model [12] is an extension of the basic LiNGAM to cases where the final assumption of the basic LiNGAM model collapses because of the presence of latent classes (hidden groups). If they exist, each latent class has a different structure. "Structure" here consists of the causal order $k(i)$, the connection strengths b_{ij}, the densities of disturbances p_i, and the means of observed variables μ_i. Therefore, structure in a class c can be described using $k^{(c)}(i)$, $b_{ij}^{(c)}$, $p_i^{(c)}$, and $\mu_i^{(c)}$ (Note that the disturbances $e_i^{(c)}$ are generated from $p_i^{(c)}$). The data within each class ($c = 1, \ldots, l$) are assumed to be generated by the basic LiNGAM. Thus, the data-generating model within class c can be described as follows:

$$x_i = \sum_{k^{(c)}(j)<k^{(c)}(i)} b_{ij}^{(c)}(x_j - \mu_j^{(c)}) + \mu_i^{(c)} + e_i^{(c)} \tag{2}$$

Moreover, the data vectors $\mathbf{x}$ are assumed to be generated by the following mixture density:

$$p(\mathbf{x}|\boldsymbol{\Theta}) = \sum_{c=1}^{l} \left\{ p\big(\mathbf{x}|\boldsymbol{\theta}^{(c)}\big) \right\} P(c) \tag{3}$$

where l is the number of classes, and $\boldsymbol{\Theta} = [\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(l)}]$, and $\boldsymbol{\theta}^{(c)}$ contain all the parameters in (2), i.e., $k^{(c)}(j), k^{(c)}(i), b_{ij}^{(c)}, \mu_j^{(c)}, \mu_i^{(c)}$, and $p_i^{(c)}(e_i^{(c)})$. Note that if only one latent class exists, the LiNGAM mixture model is equivalent to basic LiNGAM.

2.3 Existing estimation approach to LiNGAM mixture model

In this subsection, we briefly introduce the existing estimation approach to the LiNGAM mixture model. In this approach, the LiNGAM mixture model is transformed into an Independent Component Analysis (ICA) [6]) model, and an ICA algorithm (for example, [7] [9]) is applied to it [12].

The LiNGAM mixture model (Equation (2)) can be written in matrix form as follows:

$$\mathbf{x} = \mathbf{B}^{(c)}\mathbf{x} + (\mathbf{I} - \mathbf{B}^{(c)})\boldsymbol{\mu}^{(c)} + \mathbf{e}^{(c)} \tag{4}$$

where $\mathbf{B}^{(c)}$, $\boldsymbol{\mu}^{(c)}$, and $\mathbf{e}^{(c)}$ collect $b_{ij}^{(c)}$, $\mu_i^{(c)}$, and $e_i^{(c)}$, respectively, and $\mathbf{I}$ is the identity matrix. We can then obtain the following equation by solving Equation (4) for $\mathbf{x}$:

$$\mathbf{x} = \boldsymbol{\mu}^{(c)} + \mathbf{A}^{(c)}\mathbf{e}^{(c)} \tag{5}$$

where $\mathbf{A}^{(c)} = (\mathbf{I} - \mathbf{B}^{(c)})^{-1}$. We can estimate $\mathbf{B}^{(c)}$ by calculating $\mathbf{A}^{(c)}$ using the ICA algorithm and computing $\mathbf{B}^{(c)}$. However, there are two important indeterminacies that ICA cannot solve: the order and the scaling of the independent components. In order to solve them, we need to permute and normalize $\mathbf{A}^{(c)}$ appropriately before using it to compute $\mathbf{B}^{(c)}$ (See [11] for more details).

2.4 The BayesLiNGAM method

Existing estimation methods for the LiNGAM mixture model [7] [9] are affected by local optima, because of which we propose a Bayesian approach. In this subsection, we summarize the BayesLiNGAM method [5], which is an estimation method of basic LiNGAM.

In Bayesian inference, the inference with the highest posterior probability for some hypotheses is selected. In BayesLiNGAM, the hypotheses are different possible DAGs. BayesLiNGAM outputs the DAG with the highest posterior probability of all possible DAGs given the data. The posterior probabilities can be calculated by Bayes' theorem:

$$P(G_m|\mathbf{D}) = \frac{p(\mathbf{D}|G_m)P(G_m)}{p(\mathbf{D})} \tag{6}$$

Here, G_m are the different possible DAGs, $m = \{1, \ldots, N_g\}$, where N_g denotes the number of different DAGs on n variables, and $\mathbf{D}$ is the observed dataset $\mathbf{D} = \left\{\mathbf{x}^1, \ldots, \mathbf{x}^N\right\}$, N : sample size). Due to our assumption of independent and identically distributed (IID) data, $p(\mathbf{D}) = \prod_{s=1}^{N} p(\mathbf{x}^s)$. We then need to specify the likelihoods $p(\mathbf{D}|G_m)$, prior probabilities $P(G_m)$, and normalization constant $p(\mathbf{D})$ to compute the posterior probabilities.

The prior probabilities $P(G_m)$ first incorporate any domain of knowledge and prior information. If we have no knowledge or information about the G_m, all values of $P(G_m)$ are equal, i.e., $P(G_m) = \frac{1}{N_g}$.

Furthermore, $p(\mathbf{D})$ is a constant that simply normalizes posterior probabilities. Hence, we can obtain $p(\mathbf{D})$ as follows:

$$P(\mathbf{D}) = \sum_{m=1}^{N_g} p(\mathbf{D}|G_m)P(G_m) \tag{7}$$

The likelihoods $p(\mathbf{D}|G_m)$ can be computed by marginalizing over $p(\mathbf{D}|\mathbf{\Theta}, G_m)$ of $\mathbf{\Theta}$:

$$P(\mathbf{D}|G_m) = \int p(\mathbf{D}|\mathbf{\Theta}, G_m)p(\mathbf{\Theta}|G_m)d\mathbf{\Theta} \tag{8}$$

Here, $\mathbf{\Theta}$ collects all the parameters in Equation (1) (i.e., $k(j), k(i), b_{ij}, \mu_j, \mu_i$, and p_i), $p(\mathbf{D}|\mathbf{\Theta}, G_m)$ denotes the likelihood of the model (here, basic LiNGAM), and $p(\mathbf{\Theta}|G_m)$ denotes the prior distributions of $\mathbf{\Theta}$. Note that if the G_m are given, we can find the parents of x_i in the G_m and determine the causal order $k(j)$ and $k(i)$. That is, the first term of the right side in Equation (1) can be specified by the G_m. We introduce the likelihood of the model and prior distribution for $\mathbf{\Theta}$ in Section 3.

3 Proposed method

3.1 Model

We assume that the data within each class ($c = 1, \ldots, l$) are generated by the basic LiNGAM. Thus, the data-generating model in class c can be described by Equation (2), and the probability densities of the data $\mathbf{x}$ are the same as in Equation (3). We then use the Bayesian approach to estimate the LiNGAM mixture model in the same manner as BayesLiNGAM in Section 2.4.

Here, we need to parameterize the densities p_i to specify $p(\mathbf{D}|\mathbf{\Theta}, G_m)$ in Equation (8). Due to the assumption in LiNGAM that the e_i are all continuous random variables with *non-Gaussian* distributions, we use a generalized Gaussian distribution [8] that incorporates a shape parameter. The generalized Gaussian distribution is symmetric and includes Gaussian, Laplacian, continuous uniform, and several non-Gaussian distributions. The probability density function is as follows:

$$p_i(e_i) = \frac{\lambda_i \sqrt{\frac{\Gamma(3/\lambda_i)}{\Gamma(1/\lambda_i)}} \exp(-(\sqrt{\frac{\Gamma(3/\lambda_i)}{\Gamma(1/\lambda_i)}} \frac{|e_i|}{\sigma_i})^{\lambda_i})}{2\sigma_i \Gamma(1/\lambda_i)} \tag{9}$$

Here, σ_i represent the standard deviations, λ_i are the shape parameters, and $\Gamma()$ denotes the Gamma function.

Moreover, we need to transform the density of the e_i into that of the x_i. The LiNGAM mixture model (Equation (2)) can be written in matrix form as follows:

$$\begin{aligned} \mathbf{x} &= \mathbf{f}(\mathbf{e}^{(c)}) \\ &= \boldsymbol{\mu}^{(c)} + (\mathbf{I} - \mathbf{B}^{(c)})^{-1} \mathbf{e}^{(c)} \end{aligned} \tag{10}$$

where $\mathbf{f}()$ is a mapping vector, and $\mathbf{B}^{(c)}$, $\mathbf{e}^{(c)}$, and $\boldsymbol{\mu}^{(c)}$ collect $b_{ij}^{(c)}$, $e_i^{(c)}$, and $\mu_i^{(c)}$, respectively. $\mathbf{I}$ is the identity matrix.

The density of $\mathbf{x}$ within class c $\left(p(\mathbf{x}|\boldsymbol{\theta}^{(c)}, G_m)\right)$ is obtained from the density of $\mathbf{e}^{(c)}$ $\left(p_{\mathbf{e}}^{(c)}(\mathbf{e}^{(c)})\right)$ as follows [6]:

$$p(\mathbf{x}|\boldsymbol{\theta}^{(c)}, G_m) = \frac{1}{|\det J\mathbf{f}(\mathbf{f}^{-1}(\mathbf{x}))|} p_{\mathbf{e}}^{(c)}(\mathbf{f}^{-1}(\mathbf{x})) \tag{11}$$

where $J\mathbf{f}$ is a Jacobian matrix. From equations (10) and (11) with the assumption of acyclicity, $\det J\mathbf{f}(\mathbf{f}^{-1}(\mathbf{x}))$ equals to one, and we can obtain the following equation:

$$p(\mathbf{x}|\boldsymbol{\theta}^{(c)}, G_m) = p_{\mathbf{e}}^{(c)}(\mathbf{e}^{(c)}) \tag{12}$$

According to the assumption in LiNGAM that the e_i are mutually independent,

$$p_{\mathbf{e}}^{(c)}(\mathbf{e}) = \prod_{i=1}^{n} p_i^{(c)}(e_i^{(c)}). \tag{13}$$

Then, the probability density of $\mathbf{x}$ within each class c, using equations (2), (12), and (13), is given by

$$\begin{aligned}
&p(\mathbf{x}|\boldsymbol{\theta}^{(c)}, G_m) \\
&= \prod_{i=1}^{n} p_i^{(c)}(e_i^{(c)}) \\
&= \prod_{i=1}^{n} p_i^{(c)}(x_i - \mu_i^{(c)} - \sum_{k^{(c)}(j)<k^{(c)}(i)} b_{ij}^{(c)}(x_j - \mu_j^{(c)})).
\end{aligned} \tag{14}$$

We can now specify $p(\mathbf{D}|\boldsymbol{\Theta}, G_m)$ with the assumption of IID data and equations (3) and (14):

$$p(\mathbf{D}|\boldsymbol{\Theta}, G_m) = \prod_{s=1}^{N} p(\mathbf{x}^s|\boldsymbol{\Theta}, G_m) \tag{15}$$

$$\begin{aligned}
&p(\mathbf{x}|\boldsymbol{\Theta}, G_m) \\
&= \sum_{c=1}^{l} \left\{ p(\mathbf{x}|\boldsymbol{\theta}^{(c)}, G_m) \right\} P(c) \\
&= \sum_{c=1}^{l} \left\{ \prod_{i=1}^{n} p_i^{(c)}\big(x_i - \mu_i^{(c)} - \sum_{k^{(c)}(j)<k^{(c)}(i)} b_{ij}^{(c)}(x_j - \mu_j^{(c)})\big) \right\} \\
&\times P(c)
\end{aligned} \tag{16}$$

3.2 Priors for parameters and hyperparameters

We use multinomial distribution for $P(c)$ in Equation (3). The distribution represents, for N independent trials, each of which leads to one of the possible events $c = 1, \ldots, l$, the number of times that each event occurs ($\mathbf{z} = [z^{(1)}, \ldots, z^{(l)}]$). $w^{(c)}$ is the probability of the occurrence of each event. Generally, the probability mass function of a multinomial distribution is as follows:

$$P(\mathbf{z}) = \frac{N!}{\prod_{c=1}^{l} z^{(c)}!} \prod_{c=1}^{l} (w^{(c)})^{z^{(c)}} \tag{17}$$

where $w^{(c)} > 0$ and $\sum_{c=1}^{l} w^{(c)} = 1$. When we use the above distribution in Equation (3), the number of trials N is one. Then, $\mathbf{z}$ can be regarded as the indicator vector and corresponds to variable c. For example, if $\mathbf{x}$ belongs to class 2 (i.e., $c = 2$), $z^{(2)} = 1$ and the other components are 0. From the above, we can use a multinomial distribution for $P(c)$ in the following form:

$$P(c) = w^{(c)} \tag{18}$$

We use a Dirichlet distribution for the parameters of the multinomial distribution $(w^{(1)}, \ldots, w^{(l)})$ because the former is a conjugate prior for the latter

and is typically used in similar contexts. The density function of the Dirichlet distribution is as follows:

$$p(\mathbf{w}) = \frac{\Gamma(\sum_{c=1}^{l} a^{(c)})}{\prod_{c=1}^{l} \Gamma(a^{(c)})} \prod_{c=1}^{l} (w^{(c)})^{a^{(c)}-1} \tag{19}$$

where $\mathbf{w} = w^{(1)}, \ldots, w^{(l)}$ and $a^{(c)}(> 0)$ are concentration parameters. We can generate a random vector $\mathbf{w}$ by normalizing independent gamma random variables $\gamma^{(1)}, \ldots, \gamma^{(l)}$ with shape parameters $a^{(c)}$ (concentration parameters in the Dirichlet distribution) and scale parameter 1 [4]:

$$w^{(c)} = \frac{\gamma^{(c)}}{\sum_{c=1}^{l} \gamma^{(c)}} \tag{20}$$

$$p(\gamma^{(c)}) = \frac{1}{\Gamma(a^{(c)})} (\gamma^{(c)})^{a^{(c)}-1} \exp(\gamma^{(c)}) \tag{21}$$

We use Gaussian distributions for $b_{ij}^{(c)}$ and $\mu_i^{(c)}$ with zero and $\varphi^{(c)}$ as their means and v^2 and τ^2 as their variances respectively.

$$p(b_{ij}^{(c)}) = \frac{1}{\sqrt{2\pi v^2}} \exp\left\{ -\frac{(b_{ij}^{(c)})^2}{2v^2} \right\} \tag{22}$$

$$p(\mu_i^{(c)}) = \frac{1}{\sqrt{2\pi\tau^2}} \exp\left\{ -\frac{(\mu_i^{(c)} - \varphi^{(c)})^2}{2\tau^2} \right\} \tag{23}$$

We heuristically determine the value of φ . We use the result of the Gaussian mixture model estimation and an expectation-maximization (EM) algorithm [2] for φ because the greater the number of latent classes, the more difficult it is to estimate $\mu_i^{(c)}$. The EM algorithm is a method to compute maximum likelihood solutions for models with latent variables. In the algorithm, parameters and responsibilities (conditional probabilities of latent variables given the data) are updated in turn until the change in the likelihood function falls below some threshold.

We use inverse gamma distributions [3] with shape parameters α, η, and χ and scale parameters β, ζ, and ϵ for $(\sigma_i^{(c)})^2$, $\lambda_i^{(c)}$, and v^2, respectively.

$$p((\sigma_i^{(c)})^2) = \frac{\beta^\alpha}{\Gamma(\alpha)} (\sigma_i^{(c)})^{2(-\alpha-1)} \exp\left\{ \frac{-\beta}{(\sigma_i^{(c)})^2} \right\} \tag{24}$$

$$p(\lambda_i^{(c)}) = \frac{\zeta^\eta}{\Gamma(\eta)} (\lambda_i^{(c)})^{-\eta-1} \exp\left(\frac{-\zeta}{\lambda_i^{(c)}} \right) \tag{25}$$

$$p(v^2) = \frac{\epsilon^\chi}{\Gamma(\chi)} v^{2(-\chi-1)} \exp\left(\frac{-\epsilon}{v^2} \right) \tag{26}$$

We can generate inverse gamma random variables using gamma random variables. If a variable X has a gamma distribution with shape parameter α and

scale parameter β, $Y = 1/X$ has an inverse gamma distribution with shape parameter α and scale parameter $1/\beta$ [3].

We determine the value of $a^{(c)}$, α, β, η, ζ, χ, ϵ, and τ arbitrarily. From the above, we can compute $P(\mathbf{D}|G_m)$ in Equation (8) using ordinary Monte Carlo sampling [2] to compute the integral.

4 Simulation

In this simulation, we assume for simplicity that we know of the existence of a causal connection between two observed variables, and that its direction is the same for all classes. We can thus estimate which direction is true ($x_1^{(c)} \rightarrow x_2^{(c)}$ or $x_1^{(c)} \leftarrow x_2^{(c)}$). We generated 1,000 datasets under every combination of sample size ($N = 50, 100, 500$) and the number of classes ($l = 2, 4, 6$).

The data within each class were generated by the basic LiNGAM (Equation(1)), following which we mixed them. In all datasets, the true model was $x_1^{(c)} \rightarrow x_2^{(c)}$ and the connection strength was $b_{ij}^{(c)} = [-1.5, -0.5] \cap [0.5, 1.5]$. The distributions of external influences $e_i^{(c)}$ were randomly selected from the following three non-Gaussian distributions with one variance: Laplace, Uniform ($[-\sqrt{3}, \sqrt{3}]$) and Student-t (with five degrees of freedom).

With regard to the priors we did not use any prior information about G_m, and thus $P(G_m) = \frac{1}{2}$. We randomly selected the value of $a^{(c)}$ from among 3, 5, and 7. Our choice of values for the rest of the parameters was $\alpha = 3$, $\beta = 3$, $\eta = 3$, $\zeta = 3$, $\chi = 3$, $\epsilon = 3$, and $\tau = 0.5$.

We selected the number of classes l as follows. We calculated the log-marginal likelihoods of the two models for all numbers of classes and selected the class with the largest log-marginal likelihood. The maximum number of classes tested was $2\log N$. This was motivated by the fact that when the sample size N approaches infinity in a Dirichlet process mixture model, the number of classes l approaches $\log N$ [1].

We compared our method with the existing method [12] to determine the one that can correctly estimate causal directions more times given the above datasets. However, we use [9] in the existing method [12] instead of [7]. We show the simulation results in Table 1. The numbers in the tables denote the number of times each method correctly estimated causal directions. As shown in the table, our method has more correct answers than the existing method.

5 Conclusion

In this paper we proposed a new estimation approach for discovering causal structure in the presence of latent classes. In a simulation run using artificial data, our method correctly estimated more causal directions than the existing method. Our plan for future research is to evaluate our method on a wide variety of real datasets.

Table 1: Simulation results

		Sample size		
		50	100	500
Our method	$l = 2$	913	947	981
	$l = 4$	908	937	973
	$l = 6$	922	957	967
Existing method	$l = 2$	649	657	684
	$l = 4$	663	655	729
	$l = 6$	646	700	762

References

[1] C. E. Antoniak. Mixtures of dirichlet processes with applications to bayesian nonparametric problems. *The annals of statistics*, pages 1152–1174, 1974.

[2] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[3] John D Cook. Inverse gamma distribution. *online: h ttp://www. johndcook. com/inverse gamma. pdf, Tech. Rep*, 2008.

[4] L Devroye. *Non-uniform Random Variate Generation*. Springer-Verlag, 1986.

[5] P. O. Hoyer and A. Hyttinen. Bayesian discovery of linear acyclic causal models. In *Proc. 25th Conf. on Uncertainty in Artificial Intelligence (UAI2009)*, pages 240–248, 2009.

[6] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent component analysis*. Wiley, New York, 2001.

[7] M. N. H. Mollah, M. Minami, and S. Eguchi. Exploring latent structure of mixture ICA models by the minimum β-divergence method. *Neural Computation*, 18:166–190, 2006.

[8] S. Nadarajah. A generalized normal distribution. *Journal of Applied Statistics*, 32(7):685–694, 2005.

[9] J. A. Palmer, S. Makeig, K. Kreutz-Delgado, and B. D. Rao. Newton method for the ICA mixture model. In *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP2008)*, pages 1805–1808, 2008.

[10] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.

[11] S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen. A linear non-gaussian acyclic model for causal discovery. *J. Machine Learning Research*, 7:2003–2030, 2006.

[12] S. Shimizu and A. Hyvärinen. Discovery of linear non-gaussian acyclic models in the presence of latent classes. In *Proc. 14th Int. Conf. on Neural Information Processing (ICONIP2007)*, pages 752–761, 2008.

[13] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. Springer Verlag, 1993. (2nd ed. MIT Press 2000).